

Proprioceptive Inference for Dual-Arm Grasping of Bulky Objects Using RoboSimian

Matt Burkhardt¹, Sisir Karumanchi², Kyle Edelberg², Joel W. Burdick¹ and Paul Backes²

Abstract—This work demonstrates dual-arm lifting of bulky objects based on inferred object properties (center of mass (COM) location, weight, and shape) using proprioception (i.e. force torque measurements). Data-driven Bayesian models describe these quantities, which enables subsequent behaviors to depend on confidence of the learned models. Experiments were conducted using the NASA Jet Propulsion Laboratory’s (JPL) RoboSimian to lift a variety of cumbersome objects ranging in mass from 7kg to 25kg. The position of a supporting second manipulator was determined using a particle set and heuristics that were derived from inferred object properties. The supporting manipulator decreased the initial manipulator’s load and distributed the wrench load more equitably across each manipulator, for each bulky object. Knowledge of the objects came from pure proprioception (i.e. without reliance on vision or other exteroceptive sensors) throughout the experiments.

I. INTRODUCTION

Manipulating an object with two hands is a fundamental capability for many tasks in mobile manipulation, including the proper use of many tools and the lifting of bulky objects. Selecting the grasp locations on the object’s surface is thus a prerequisite for task completion. In this paper, our task is to lift and support bulky objects whose mass and size (e.g. ~ 50 lbs and ~ 1 m maximum weight and extent) could either overload the capacity of a single manipulator or significantly degrade mobility performance after the payload is lifted with a single manipulator. Dual arm grasping can benefit both these scenarios by spreading load more evenly across the robot-payload system. The challenge is to identify two grasp locations on the object that minimize the wrench loads experienced at each manipulator. Knowledge of the object’s shape, center of mass (COM) location, and weight are not known *a priori*. However, these properties strongly affect manipulability and confidence in their values is necessary for robust manipulation in the field. To solve this task, we present an algorithm that represents the object’s properties using Bayesian methods. The object’s properties are queried proprioceptively and we show experimentally that dual arm manipulation is enabled for a wide range of objects using



Fig. 1: RoboSimian Probing the Properties of a Bulky Chair

these measurements alone³. Candidate grasp locations are represented using a particle set. A heuristic that depends on confidence in the object’s properties is associated with each particle. When the heuristic of the most likely particle passes a user-defined threshold, a grasp is attempted. If no particle passes this threshold, then the algorithm chooses a property of the object to explore proprioceptively and the process is repeated.

The platform used in this paper is the Jet Propulsion Laboratory’s (JPL) RoboSimian (see Fig.1). RoboSimian is an agile quadruped designed to operate in hazardous environments. Each RoboSimian limb is composed of three elbows and seven total degrees of freedom (DOF). On each limb’s distal end is a Cam-Hand that contains four fingers and three additional DOF. A set of differential drive wheels are mounted to RoboSimian’s back and two passive wheels to its rear limbs to enable fast rolling across level ground when appropriate. RoboSimian’s flexible design enables it to alternate between extreme terrain mobility (using its four limbs) and object manipulation (using any two of its limbs). These capabilities were showcased in the DARPA Robotics Challenge (DRC) [1], [2]. In our experiments, objects are grasped with the front limb’s Cam-Hands while RoboSimian is in a sitting pose, which is shown in Fig. 1.

II. BACKGROUND, RELATED WORK & CONTRIBUTIONS

This work is motivated by scenarios in which noisy or forceful interactions are necessary to manipulate unknown objects in cluttered environments. Interactions tend to be

¹ California Institute of Technology, Mechanical and Civil Engineering Department mburkhar@caltech.edu

²NASA Jet Propulsion Laboratory (JPL), Robotic Manipulation and Sampling Group Sisir.B.Karumanchi@jpl.nasa.gov

The research described in this paper was carried out at the Jet Propulsion Laboratory, California Institute of Technology, under a contract with the National Aeronautics and Space Administration. This research was sponsored by the Army Research Laboratory and was accomplished under Cooperative Agreement Number W911NF-10-2-0016. ©2017 California Institute of Technology. Government sponsorship acknowledged. In addition, this work was partially supported by a NASA Space Technology Research Fellowship (NSTRF).

³This algorithm is sensor agnostic, e.g., range data from depth cameras or stereo vision can be readily incorporated

challenging due to unknown inertial properties of the target object, collisions with surrounding clutter and also occlusions that reduce effectiveness of exteroception driven behaviors. For these reasons, we focus on leveraging probabilistic inference methods with pure proprioception driven adaptation to tackle and reason about noisy interactions. The target’s properties are inferred using Bayesian methods, which are probabilistic estimates that update a prior assumption on the property’s value based on Bayes’ theorem (see [3] for a review).

Many works have focused on modeling an unknown object’s surface using a specific type of Bayesian method known as a Gaussian Process (GP) to enable object identification or grasp planning. A GP can be thought of as a defining a Gaussian distribution over functions,

$$f(\mathbf{x}) \sim \mathcal{GP}(m(\mathbf{x}), k(\mathbf{x}, \mathbf{x}'))$$

where $\mathbf{x} \in \mathbb{R}^d$ is a feature, $m(\mathbf{x}) = \mathbb{E}[f(\mathbf{x})]$ is the mean function and,

$$k(\mathbf{x}, \mathbf{x}') = \mathbb{E}[f(\mathbf{x} - m(\mathbf{x}))(f(\mathbf{x}') - m(\mathbf{x}'))]$$

is the covariance function [4]. The value of the function, $f(\mathbf{x}) \in \mathbb{R}$ is distributed as a Gaussian and is completely determined by $\{m(\mathbf{x}), k(\mathbf{x}, \mathbf{x}')\}$, which are selected by the user. The covariance function reflects the correlation between nearby data points $(\mathbf{x}, \mathbf{x}')$, i.e. the smoothness of the data.

GPs are data-driven processes. Given a set of training data $\mathcal{D} = \{(\mathbf{x}_i, y_i)\}_{i=1}^n$, one may obtain the posterior distribution over the function’s value at a new data point x^* . The posterior distribution is given by $f(\mathbf{x}^*) \sim \mathcal{N}(\bar{f}_*, \Sigma_*)$ where,

$$\bar{f}_* = \mathbf{k}_*^T (K + \sigma_n^2 I)^{-1} \mathbf{y} \quad (1)$$

$$\Sigma_* = k(\mathbf{x}_*, \mathbf{x}_*) - \mathbf{k}_*^T (K + \sigma_n^2 I)^{-1} \mathbf{k}_* \quad (2)$$

where $\mathbf{k}_* = \{k(\mathbf{x}_i, \mathbf{x}^*)\}_{i=1}^n$, K is an $n \times n$ matrix whose i^{th} column is \mathbf{x}_i , $\mathbf{y} = \{y_i\}_{i=1}^n$, and σ_n^2 is the signal variance [4].

Specifically GPs have been effective in representing implicit surfaces as a means to capture shape information of unknown objects. In [5], a probabilistic signed distance function (SDF) was used to represent an unknown object’s surface and uncertainty from this model was used to choose precision grasping locations. In [6], an GP implicit-surface (GPIS) represented the object’s surface. The posterior’s mean function was then used to create a potential field that guided subsequent reaching motions. Mahler et. al [7] similarly used GPIS to represent a novel object’s surface, which was then used to orient a parallel-jaw gripper such that the resulting grip would possess a high probability of force closure. Choosing to represent an object’s surface probabilistically allows subsequent tasks to be conditioned on the uncertainty in the surface model, which may trigger exploitation of current model or further exploration (e.g. using haptics to probe uncertain areas of an object as in [8]). Although vision is the predominant sense used to learn about an object’s properties, several researchers have used proprioception to learn about novel objects without the use of computer vision.

In [9], a manipulator interacted with an unknown object according to a set of primitives (e.g. push, shake, lift) while recording joint torques and auditory sounds. The collected data produced features that were used to classify the object. Similarly, classification in [10] followed from manipulating an object according to an “excitation trajectory,” which allowed the object’s mass, COM location, and inertia tensor to be used as features. Furthermore, classification features in [11] were based on finger joint angles that were recorded while grasping an object. These joint angles conveyed a “signature” of the object’s shape. Although most of this research has focused on object identification, one can imagine learning these object properties to enable other behaviors like lifting a bulky object.

Thus, the key contributions of this work include (i) a probabilistic representation of an object’s shape, COM location, and weight based on pure proprioception and (ii) the use of this representation to identify grasping locations to permit the dual arm support of bulky objects. Compared to previous work in this area, there are two key reasons for the significance and novelty of this work for Robotics community. First, objects properties are inferred not for the purpose of recognition but to generate reactive behaviors on the fly. The goal here is develop the simplest behavior that can work with extremely sparse information. Second, experimental tests demonstrate a adaptive behavior on heavy and bulky objects (10kg to 25kg). To the knowledge of the authors, experimental demonstrations beyond 10kg payloads have not been shown before with human-scale mobile manipulation robots.

The remainder of the paper is organized as follows. In Section III, the probabilistic model describing an object’s COM location and mass are described. Section IV introduces the probabilistic model of an object’s shape. Section V presents our method of selecting a grasping location using particles, followed by future work detailed in Section VI. Experimental results can be found throughout the remaining sections.

III. OBJECT CENTER OF MASS (COM) LOCATION & WEIGHT

The object’s COM location and weight are inferred using a six DOF force-torque (FT) sensor mounted between the distal end of RoboSimian’s limb and the Cam-Hand.

A. Estimating COM Position

Lifting an object with one hand and recording the FT is sufficient to determine two DOF of the COM’s position. The DOF pertaining to the COM’s position along the direction of gravity is unobservable using a single measurement. Thus, the object must be rotated to ascertain the full three DOF defining the COM position. For simplicity, we initially assume that one manipulator can lift and rotate the object. All estimates of the object’s COM are calculated with respect to a frame EE , fixed to this first manipulator.

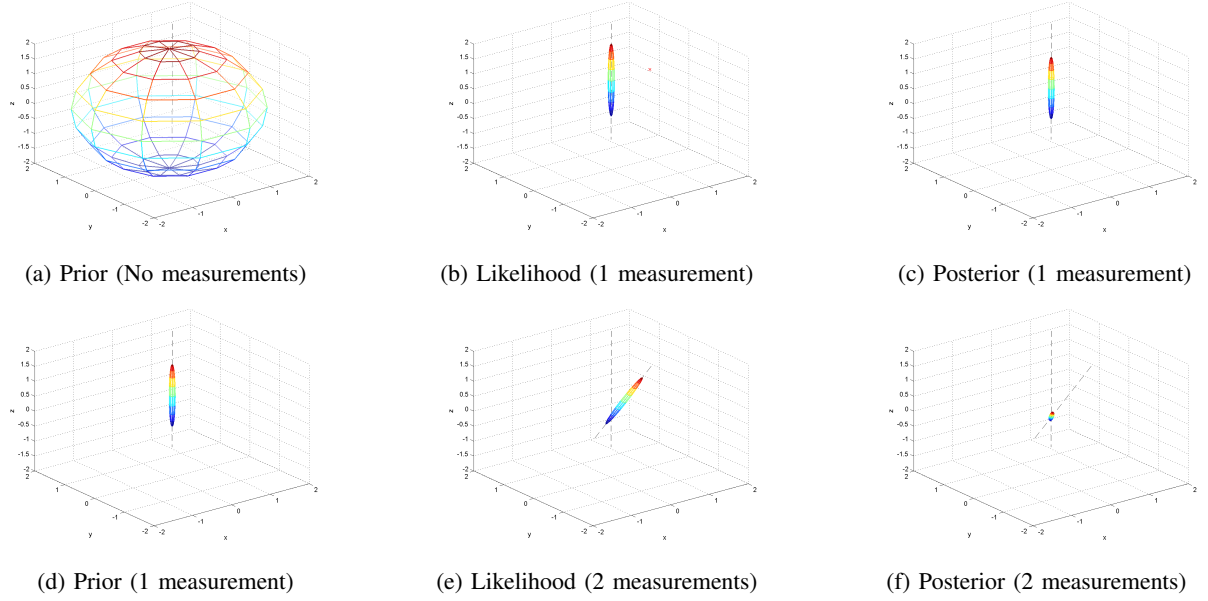





Fig. 3: Illustrative example of COM Inference (dotted lines show the local direction of gravity in the EE frame)



(a) Pallet (14.2 kg)



(b) User Display with COM (dark red ellipsoid) overlay'ed with a stereo point cloud



(c) COM Magnitude, $\|\mathbf{r}_s\|$

Fig. 4: Probing the COM location of a pallet

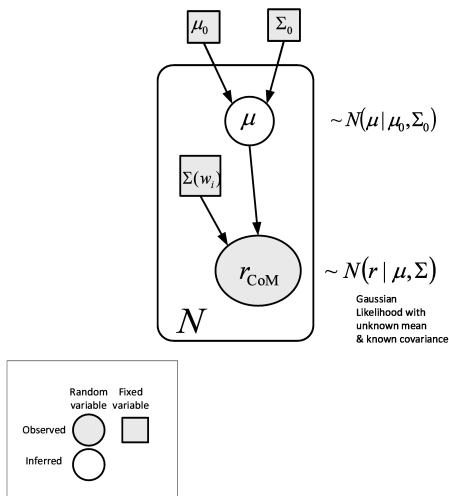


Fig. 2: Graphical representation of the COM's probabilistic model.

Estimation of the object's COM location follows the block diagram shown in Fig. 2. The COM position expressed in EE is given by,

$$\mathbf{r}_{COM} = \mathbf{r}_s + \epsilon \mathbf{1}$$

where $\mathbf{r}_s \in \mathbb{R}^3$ is the static COM position, $\epsilon \in \mathbb{R}$ is random noise and $\mathbf{1} = \{1\}_{i=1}^3$. Typically, the measurement model for \mathbf{r}_s would include two successive rotations of the manipulator to fully constrain the COM's position. However, a single FT measurement without subsequent rotation can be used if the measurement model is heteroscedastic. A heteroscedastic model is one in which the noise is a smooth function of the inputs [12]. Let \mathbf{r}_s be governed by a multidimensional Gaussian, $\mathbf{r}_s \sim \mathcal{N}(\boldsymbol{\mu}, \Sigma(w_{ee}))$ with an unknown mean $\boldsymbol{\mu}$ and an input-dependent variance $\Sigma(w_{ee})$, where $w_{ee} = \{f_{ee}, \tau_{ee}\}$ is a single wrench measurement expressed in frame EE . The variance of $\Sigma(w_{ee})$ is a constant small-value in the directions orthogonal to gravity since a single wrench measurement constrains these DOF. Along the gravity direction, the vari-

ance is a constant large value such that the $\Sigma(w_{ee})$ resembles an elongated ellipsoid. Measurements of the COM follow from the wrench,

$$\boldsymbol{\mu}_m = \frac{f_{ee} \times \boldsymbol{\tau}}{\|f_{ee}\|^2}$$

This measurement is physically equivalent to a projection of the true COM location onto a plane orthogonal to the local direction of gravity in EE . Although this measurement captures two DOF of the COM, the full three DOF can be achieved with rotation of the object and the heteroscedastic model. The noise is distributed as $\epsilon \sim \mathcal{N}(0, \lambda^{-1})$, where λ is a constant precision¹. Then,

$$\mathbf{r}_{COM} \sim \mathbf{N}(\boldsymbol{\mu}, \Sigma(w_{ee}) + \lambda^{-1}I)$$

A distribution governing $\boldsymbol{\mu}$ can be obtained by setting a multidimensional Gaussian prior $\mathbf{N}(\boldsymbol{\mu}_0, \Lambda_0)$. Together with the likelihood $\mathbf{N}(\boldsymbol{\mu}, \Sigma(w_{ee}))$, the posterior governing $\boldsymbol{\mu}$ can be obtained in closed-form [3] as $\boldsymbol{\mu} \sim \mathbf{N}(\boldsymbol{\mu}_n, \Lambda_n)$ where,

$$\begin{aligned} \boldsymbol{\mu}_n &= (\Lambda_0^{-1} + n\Sigma^{-1})^{-1}(\Lambda_0^{-1}\boldsymbol{\mu}_0 + n\Sigma^{-1}\bar{\boldsymbol{\mu}}) \\ \Lambda_n^{-1} &= \Lambda_0^{-1} + n\Sigma^{-1} \end{aligned}$$

where $\bar{\boldsymbol{\mu}}$ is the average of the measurements $\boldsymbol{\mu}_m$. This posterior is then used to reason about \mathbf{r}_{COM} , as shown in Fig. 3. In Fig. 3b, the elongated form of $\Sigma(w_{ee})$ is evident after a single measurement from the FT. Note that the true COM may lie anywhere on this line. In Fig 3e, the wrist has been rotated and an additional measurement has been taken leading to another elongated likelihood for the COM but in another direction due to the hand's motion. The posterior in Fig. 3f after two measurements is localized on the true location of the COM and its sensitivity is a function of the angle through which the object is rotated. It is interesting to observe that the heteroscedastic model adopted here enables us to deduce the three DOFs of the true COM location despite the two DOF measurements, $\boldsymbol{\mu}_m$, that are possible with the FT sensor.

Figure 4 demonstrates the COM probe algorithm for a pallet. Figure 4a shows a view from RoboSimian's camera as the CamHand grasps the pallet. Figure 4b shows the user's interface after the COM probe, where certainty in COM position is denoted by a one standard-deviation ellipsoid shaded in maroon. Figure 4c demonstrates the convergence of $\|\mathbf{r}_s\|$ as the object is rotated 30° from the orientation shown in Fig. 4a. It is evident that $\|\mathbf{r}_s\|$ stabilizes after three FT measurements taken at different object orientations, as can be seen from the blue line in Fig. 4c. The red line of Fig. 4c shows the COM position responding to an external, unexpected wrench (the object was kicked during FT acquisition). As can be expected, the variance on $\|\mathbf{r}_s\|$ rose due to the disturbance and settled to a stable value after more undisturbed data was collected.

¹A stochastic precision model has been developed to enable the variance to inflate when disturbance forces or collisions perturb the object during the COM probe.

B. Estimating Weight

The weight of the object, ω , is similarly modeled as a stochastic variable. Weight measurements are provided by the norm of the measured FT force, $\|f_{ee}\|$. The likelihood is modeled by a simple Gaussian with unknown mean and precision, $\mathcal{N}(\mu, \lambda^{-1})$. The conjugate prior to this likelihood is the normal-Gamma distribution, $\mathcal{NG}(\mu_0, \kappa_0, \alpha_0, \beta_0)$ (see [13] for details). The posterior is then a normal-Gamma distribution, where $\mathcal{NG}(\mu, \lambda|\mu_0, \kappa_0, \alpha_0, \beta_0)$ and,

$$\begin{aligned} \mu_n &= \frac{\kappa_0\mu_0 + n\bar{\omega}}{\kappa_0 + n} \\ \kappa_n &= \kappa_0 + n \\ \alpha_n &= \alpha_0 + \frac{n}{2} \\ \beta_n &= \beta_0 + \frac{1}{2} \sum_{i=1}^n (m_i - \bar{\omega})^2 + \frac{\kappa_0 n (\bar{\omega} - \mu_0)^2}{2(\kappa_0 + n)} \end{aligned}$$

where $\bar{\omega}$ is the mean of all n weight measurements. The normal-Gamma posterior then fully characterizes the distribution of ω and in practice μ_n is adopted as the object's weight.

IV. OBJECT SHAPE

The object's shape is represented using GP implicit surfaces (GPIS). An implicit surface is defined by a smooth function $\phi(\mathbf{x}) = 0$, where $\mathbf{x} \in \mathbb{R}^3$ represents a position in ambient space. We assume that the target is a single, closed object. Let \mathcal{O} denote the points defining the object's surface. Similar to [6], we have adopted the convention,

$$\phi(\mathbf{x}) \text{ is } \begin{cases} > 0, & \mathbf{x} \text{ outside } \mathcal{O} \\ < 0, & \mathbf{x} \text{ inside } \mathcal{O} \\ = 0, & \mathbf{x} \in \mathcal{O} \end{cases}$$

The GP is used to predict the function $\phi : \mathbb{R}^3 \rightarrow \mathbb{R}$. For each ambient position \mathbf{x} , the GP will return a posterior mean $\mu(\mathbf{x})$ (from Eq. 1) and covariance $\sigma(\mathbf{x})$ (from Eq. 2) on the function $\phi(\mathbf{x})$. Normals to the surface can also be calculated by taking the gradient of Eq. 1 with respect to \mathbf{x} .

We chose a thin-plate (TP) covariance function [6] for our GP model. This covariance function is a popular choice to model implicit surfaces as it was shown to have the desired behavior at boundaries [6]. In addition, this function is desirable as it only has one hyper-parameter and given that we want to operate with unknown objects it is hard to get access to a good training set. This function is defined as follows:

$$k_{TP}(\mathbf{x}, \mathbf{x}') = \frac{1}{12}(2d^3 - 3Rd^2 + R^3) \quad (3)$$

where $d \triangleq \|\mathbf{x} - \mathbf{x}'\|$ and R is a constant that physically corresponds to the maximum anticipated value between any two points in the feature space. We trained the hyper-parameter R of the thin-plate covariance function by fitting to a representative shape like a cube as shown in Fig. 5.

A set of exterior boundary points, \mathcal{S} , are added as training data to condition the GP. These points are equally distributed about a sphere of constant radius and $\phi(\mathcal{S}) = +\sigma$, where σ is a large positive value. When a manipulator probes the object, the contact point \mathbf{x}_{con} in frame EE is added to the GP's training data with a value of zero. In addition, the contact normal, \mathbf{n}_{con} , is estimated using the measured force from the FT. Using this normal, we add two additional points $\phi(\mathbf{x}_{con} \pm \varepsilon \mathbf{n}) = \pm \varepsilon$ to the training data, where ε is a small constant parameter. This proved to be a simple way to implement surface normal measurements into the GP. When a manipulator moves towards the object and fails to make contact, negative information (i.e. the lack of contact) is also used to update the shape GP. The result of the shape GP on a grasped stool after some probing is shown in Fig. 6.




Fig. 5: Thin-plate covariance function's parameter R was chosen by training on a known object like a cube. The colored mesh denotes the zero level-set of the implicit surface and the warm colors denote uncertainty in the shape.




Fig. 6: Result of shape GP on a grasped stool after ~ 14 points have been probed around the stool's exterior. A stereo point cloud of the true object is shown along with the zero level-set of the implicit surface.

V. CHOOSING GRASP LOCATION & EXPERIMENTS

We assume that one manipulator initially grasps and lifts the object. Once the object's shape, COM, and weight have been determined, the task is to select the position of a supporting manipulator that we assume contacts the object at a point without friction. To facilitate this, a set of particles is

distributed uniformly throughout the workspace. Each particle represents the location of the supporting manipulator, \mathbf{x} . The force imparted by the supporting manipulator is assumed to be equal to half the weight of the object. Heuristics are associated with each particle in the set, corresponding to the likelihood of lifting the bulky object. The heuristics for success come from (i) confidence that the contact location is on the object's surface, (ii) equality of the wrench norms between the two manipulators and (iii) confidence in the object's measured weight and COM location.

A. Shape Heuristic

The supporting manipulator's contact point must be confidently located on the object's surface, $\mathbf{x} \in \mathcal{O}$. The shape GP provides the posterior mean value of the implicit function $\mu(\mathbf{x})$ and the covariance $\sigma(\mathbf{x})$. A heuristic rewarding particles that are confidently located on the object's surface can be written as,

$$h_{shape} = \frac{1}{2}e^{-C_1\mu(\mathbf{x})^2} + \frac{1}{2}e^{-C_2\sigma(\mathbf{x})^2}$$

where $h_{shape} \in [0, 1]$ and (C_1, C_2) are positive. This heuristic rewards particles that are located on the surface $\mu(\mathbf{x}) = 0$ with low covariances $\sigma(\mathbf{x})$.

B. Wrench Equality Heuristic

The objective is to pick the supporting manipulator's contact position to minimize the wrench norm experienced at each manipulator. Further, it is desirable that the wrench norms at each manipulator be equal. Let the hand frame of the initially-contacting manipulator be denoted EE_1 while the second hand frame is denoted EE_2 . Assuming that the second manipulator contacts the surface at a point with no friction, static equilibrium of the object is written as,

$$w_r^1 = -w_g^1 - w_2^1(\mathbf{d}_2)$$

where w_g^1 is the wrench due to gravity, w_2^1 is the wrench due to the second manipulator, and w_r^1 is the remainder wrench required for static equilibrium, all expressed in frame EE_1 . The wrench acting on EE_1 is w_r^1 . Again using static-equilibrium, we can express the wrench felt at EE_2 by,

$$w_r^2 = Ad_{g_{12}}^T w_r^1$$

where $g_{12} \in SE(3)$ is the rigid transformation between the frames EE_1 and EE_2 . For a given wrench $w = (f, \tau) \in \mathbb{R}^6$, we assume a wrench norm of $\|w\| = \|f\| + \|\tau\|$. Then, we can form a heuristic over the particle set that enforces wrench equality and minimization as,

$$h_{we} = \frac{1}{2}e^{-C_1\left(1 - \frac{\|w_r^1\|}{\|w_r^2\|}\right)^2} + \frac{1}{2}e^{-C_2\left(1 - \frac{\|w_r^2\|}{\|w_{r,min}^2\|}\right)^2}$$

where $w_{r,min}^2$ is the smallest wrench norm of any particle in the set, and (C_1, C_2) are positive. This heuristic lies in the range $[0, 1]$ with an optimal particle having unity weight.

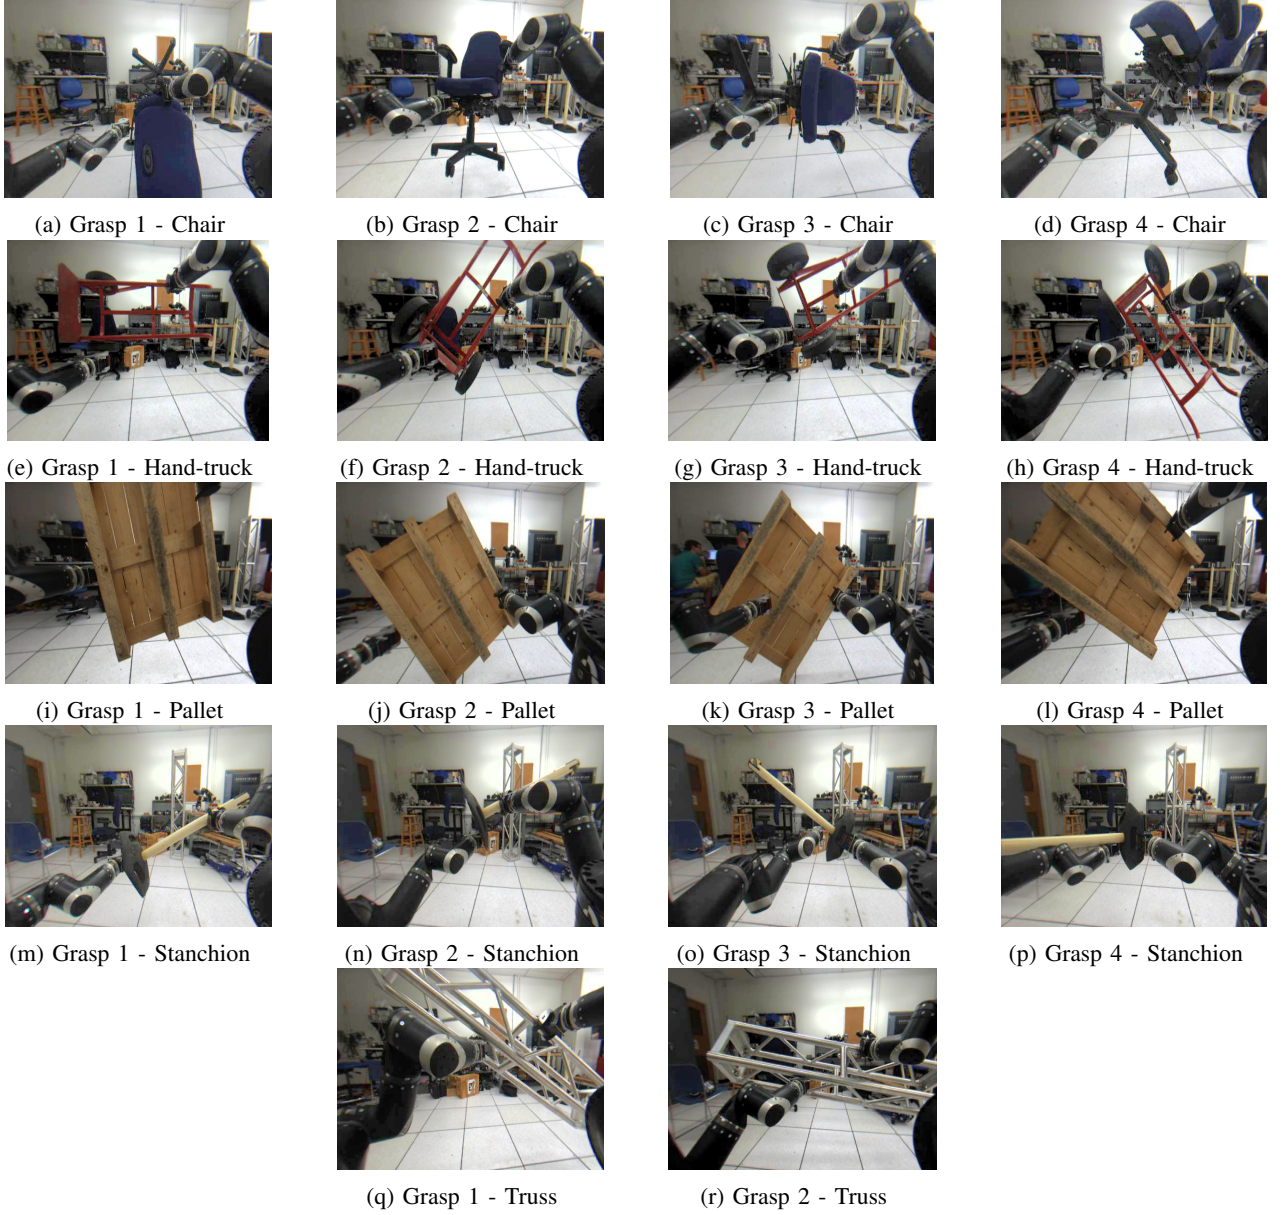


Fig. 7: Ending postures of an auto-adaptive dual arm behavior given different operator-specified initial grasp locations. The results are shown for a chair (25 kg), hand-truck(14.6 kg), pallet(14.2 kg), stanchion (7.6 kg) and truss (17.2 kg).

C. COM & Weight Heuristic

Lifting a bulky object requires confidence in the object's COM location and weight. Uncertainty in the COM location or weight would jeopardize a lift no matter where the supporting hand is placed. Therefore, a heuristic based on the certainty of the weight and COM should penalize the particles uniformly. From Sec. III, the posterior covariance on the COM's position is given by Λ and the variance on the weight follows from λ . The following heuristic penalizes each of the particles uniformly,

$$h_m = \frac{1}{2}e^{-C_1\text{Tr}(\Lambda)} + \frac{1}{2}e^{-C_2\sigma}$$

where again (C_1, C_2) are positive and $h_m \in [0, 1]$ with favorable particles near unity.

D. Updating Particle Set

The sum of the heuristics on each particle are used as weights during a re-sampling step. A Low-Variance Sampler (LVS) re-samples the particle set, with replacement, according to the weights associated to each particle [14]. This has the effect of populating the set with particles of high heuristic value, i.e. high likelihood of a stable lift, while still maintaining particle diversity. Since re-sampling occurs with replacement, a single particle may be included in the set multiple times. The relative number of times a particle is included in the set may be interpreted as a likelihood of a stable grasp.

E. Experiments

Our experiments follow the state machine shown in Fig. 8. If the maximum particle weight is below a user-defined threshold h_{th} , then a grasp is not attempted due to a lack of confidence in the object's properties. Determining the object property with the greatest uncertainty follows by evaluating the relative fraction ($\alpha = 1 - \frac{h_{sp}}{h_m}$) over the particle set. For particles in the set where $\alpha < 0$, this implies $h_m < h_{sp}$, signifying that knowledge of the COM position or weight must improve. This triggers an additional COM probe whereby the object is rotated as in Section III-A. Likewise, for particles where $\alpha > 0$, this implies $h_{sp} < h_m$, which signifies that greater knowledge of object shape is desired. This triggers the second manipulator to probe the object near the most uncertain portions of the GPIS.

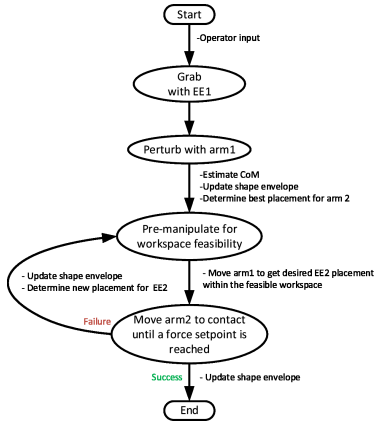


Fig. 8: A state machine illustrating a basic auto-adaptive behavior for second end effector placement given first end effector placement via operator input.

Experiments were conducted on a variety of bulky objects with masses ranging from 7 to 25 kilograms, as shown in Fig. 7. In this figure, RoboSimian's right hand first lifts the object and estimates its COM and weight via a slight rotation of the wrist. At this point, knowledge of the object's shape is limited to the single point of contact made by the right hand. Using knowledge of the object's COM, the left manipulator moves to contact such that the COM is situated symmetrically between both manipulators. Force-torque measurements were made at the distal end of RoboSimian's right and left hand before and after the second manipulator moved to support the object. The wrench magnitudes were calculated for each object and trial, as shown in Fig. 9. Wrench norms have been rescaled such that the initial wrench norm on the right hand is unity. For the desired performance, one would require the wrench norms for both manipulators to settle at 0.5 indicating symmetric load balancing.

These plots show load transfer performance after one successful contact with the second end effector. The demonstrated trends show promise in terms of monitoring load transfer and adapting. Table I quantifies the wrench norm trends in terms of percent load reduction on end effectors. It also describes and documents observed failures. The

Object	No. of runs	% reduction in EE1 wrench norm	% increase in EE2 wrench norm	No. of failures	Comments
Chair	4	9.1(max), -0.1(min)	38.4(max), -0.7(min)	0/4	Insufficient grip on EE1, observed object slip in hand.
Handtruck	5	64.3(max), 24.8(min)	81.9(max), 46.9(min)	1/5	Insufficient grip on EE1.
Pallet	4	15.0(max), -8.9(min)	51.9(max), 37.5(min)	0/4	
Stanchion	7	13.4(max), -2.3(min)	49.0(max), 20.0(min)	3/7	Workspace constraints
Truss	5	22.4(max), 21.1(min)	19.4(max), 15.0(min)	3/5	Three failures; object-robot collision, workspace constraint, and insufficient grip.

TABLE I: Results: Failures and % change in load

hand-truck and the truss demonstrated the best symmetric load transfer after one adaptive contact behavior. The chair was the heaviest object showed minimal reduction in load of the first despite significant increase in the second end effector load (indicating build up of internal squeeze forces). Similarly, the stanchion and the pallet also showed cases where the first end effector load reduction was minimal. These observations indicate the need for further iterative information gathering and adaptation. In these preliminary experiments, our goal was to determine the simplest behavior that would work across different objects and different initial grasps without failures. In future work, we will expand the behavior to iterate until a better load transfer is achieved across end effectors. We also believe relieving internal grip forces in conjunction with adjusting end effector placement is an important future area of improvement.

Three main challenges were encountered during these tests. First, the size of the objects relative to the workspace led to collisions between RoboSimian and the target object; this was particularly noteworthy when manipulating the truss, which prohibited the truss from being handled at its ends and terminated two of the trials. Secondly, the grasps of RoboSimian's Cam-Hands had difficulty restraining the in-hand rotation of objects with low compliance. Objects that slipped in the hands included the chair, hand-truck, and truss with noticeable results on the wrench norms of Fig. 9. Third, the supporting hand occasionally moved to contact the object in a way that increased the reactive wrench on the grasping manipulator. This occurred most noticeably on the stanchion, and was due to the limited shape information that was incorporated into these tests.

VI. CONCLUSION

Our experiments have demonstrated that lifting bulky objects with two manipulators can be facilitated by taking proprioceptive measurements of the object's COM, weight, and shape. Experiments using RoboSimian to lift bulky objects up to 25kg showed that basic knowledge of the object's COM and shape could be used to position a second manipulator to balance the wrenches experienced at either hand. This basic capability is a prerequisite task to more complicated behaviors such as component assembly or the rearrangement of bulky objects in a room. Furthermore, object properties are inferred probabilistically using data-driven methods such that subsequent actions can be conditioned on confidence in our knowledge of the object.

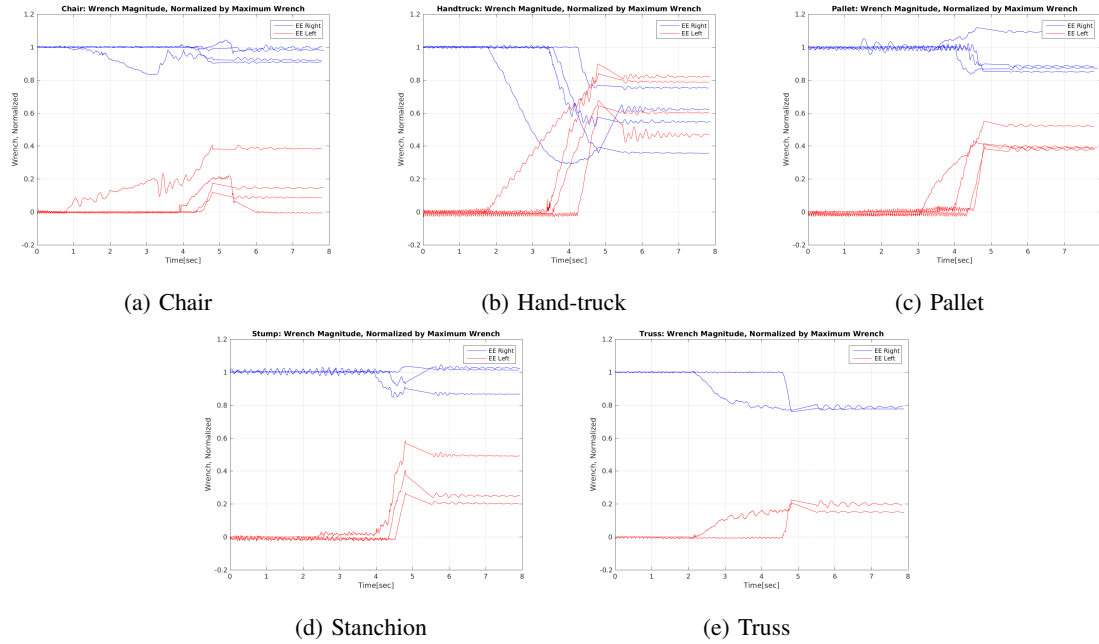


Fig. 9: Wrench norms experienced by right (blue) and left (red) manipulators. Wrench norms have been rescaled such that the initial wrench norm on the right hand is unity. For ideal performance, one would expect the wrench norms for both manipulators to settle at 0.5 indicating equal load balancing.

Future research could be productively focused on three key areas. First, studying the relationship between the evolving shape of the object and the robot's workspace could eliminate object-robot collisions. A path-planning algorithm that considers the object's shape will likely be required to place the second manipulator in a favorable position to explore difficult-to-reach areas of the object, like the backside of the pallet in Fig. 7. Second, implementing a dynamic particle set could improve the efficiency of searching for the best end-effector support point. Currently, our method resamples a fixed particle set where each particle represents an end-effector position. Using a dynamic set, perhaps based on Particle Swarm Optimization (PSO) strategies [15], would enable a smaller number of particles to explore a larger-dimensional space. This would increase the efficiency of the search as well as allow inclusion of forces into the particle definition. Third, future work could incorporate vision-derived information into the object's surface model. This would dramatically reduce the number of probing touches required provided that the vision system could operate robustly in the field.

REFERENCES

- [1] P. Hebert, M. Bajracharya, J. Ma, N. Hudson, A. Aydemir, J. Reid, C. Bergh, J. Borders, M. Frost, M. Hagman, *et al.*, "Mobile manipulation and mobility as manipulation, design and algorithms of RoboSimian," *Journal of Field Robotics*, vol. 32, no. 2, pp. 255–274, 2015.
- [2] S. Karumanchi, K. Edelberg, I. Baldwin, J. Nash, J. Reid, C. Bergh, J. Leichty, K. Carpenter, M. Shekels, M. Gildner, *et al.*, "Team RoboSimian: Semi-autonomous Mobile Manipulation at the 2015 DARPA Robotics Challenge Finals," *Journal of Field Robotics*, vol. 34, no. 2, pp. 305–332, 2017.
- [3] A. Gelman, J. B. Carlin, H. S. Stern, D. B. Dunson, A. Vehtari, and D. B. Rubin, *Bayesian data analysis*. CRC press Boca Raton, FL, 2014, vol. 2.
- [4] C. E. Rasmussen and C. K. Williams, *Gaussian processes for machine learning*. MIT press Cambridge, 2006, vol. 1.
- [5] D. Chen, V. Dietrich, and G. von Wichert, "Precision grasping based on probabilistic models of unknown objects," in *Robotics and Automation (ICRA), 2016 IEEE International Conference on*. IEEE, 2016, pp. 2044–2051.
- [6] S. Dragiev, M. Toussaint, and M. Gienger, "Gaussian process implicit surfaces for shape estimation and grasping," in *Robotics and Automation (ICRA), 2011 IEEE International Conference on*. IEEE, 2011, pp. 2845–2850.
- [7] J. Mahler, S. Patil, B. Kehoe, J. Van Den Berg, M. Ciocarlie, P. Abbeel, and K. Goldberg, "Gp-gpis-opt: Grasp planning with shape uncertainty using gaussian process implicit surfaces and sequential convex programming," in *Robotics and Automation (ICRA), 2015 IEEE International Conference on*. IEEE, 2015, pp. 4919–4926.
- [8] S. Dragiev, M. Toussaint, and M. Gienger, "Uncertainty aware grasping and tactile exploration," in *Robotics and Automation (ICRA), 2013 IEEE International Conference on*. IEEE, 2013, pp. 113–119.
- [9] J. Sinapov, T. Bergquist, C. Schenck, U. Ohiri, S. Griffith, and A. Stoytchev, "Interactive object recognition using proprioceptive and auditory feedback," *The International Journal of Robotics Research*, vol. 30, no. 10, pp. 1250–1262, 2011.
- [10] D. Kubus, T. Kroger, and F. M. Wahl, "On-line rigid object recognition and pose estimation based on inertial parameters," in *Intelligent Robots and Systems, 2007. IROS 2007. IEEE/RSJ International Conference on*. IEEE, 2007, pp. 1402–1408.
- [11] A. Vásquez, Z. Kappasov, and V. Perdureau, "In-hand object shape identification using invariant proprioceptive signatures," in *Intelligent Robots and Systems (IROS), 2016 IEEE/RSJ International Conference on*. IEEE, 2016, pp. 965–970.
- [12] P. W. Goldberg, C. K. Williams, and C. M. Bishop, "Regression with input-dependent noise: A gaussian process treatment," in *Advances in neural information processing systems*, 1998, pp. 493–499.
- [13] K. P. Murphy, "Conjugate Bayesian analysis of the Gaussian distribution," *def*, vol. 1, no. 2 σ^2 , p. 16, 2007.
- [14] S. Thrun, W. Burgard, and D. Fox, *Probabilistic robotics*. MIT press, 2005.
- [15] F. Marini and B. Walczak, "Particle swarm optimization (PSO). A tutorial," *Chemometrics and Intelligent Laboratory Systems*, vol. 149, pp. 153–165, 2015.